

# EXHIBIT R

## Expanding access to safer AI with Amazon

Sep 25, 2023 • 3 min read



Today, we're announcing that Amazon will invest up to \$4 billion in Anthropic. The agreement is part of a broader collaboration to develop the most reliable and high-performing foundation models in the industry. Our frontier safety research and products, together with Amazon Web Services' (AWS) expertise in running secure, reliable infrastructure, will make Anthropic's safe and steerable AI widely accessible to AWS customers.

AWS will become Anthropic's primary cloud provider for mission critical workloads, providing our team with access to leading compute infrastructure in the form of AWS Trainium and Inferentia chips, which will be used in addition to existing solutions for model training and deployment. Together, we'll combine our respective expertise to collaborate on the development of future Trainium and Inferentia technology.

Based on significant AWS customer demand for Claude, we're also expanding our support of [Amazon Bedrock](#). This will include secure model customization and fine-tuning on the service to enable enterprises to optimize Claude's performance with their expert knowledge, while limiting the potential for harmful outcomes.

Amazon developers and engineers will be able to build on top of our state-of-the-art models via Amazon Bedrock. This will enable them to incorporate generative AI capabilities into their work, enhance existing applications, and create net-new customer experiences across Amazon's businesses.

Organizations will be able to use [Claude 2](#) for a wide range of tasks, from sophisticated dialogue and creative content generation to complex reasoning and detailed instruction. Our industry-leading 100,000 token context window will help them securely process extensive amounts of information—including technical, domain-specific documents for use cases across finance, legal, coding, and more.

Enterprises across many industries are already building with Anthropic models on Amazon Bedrock. **LexisNexis Legal & Professional**, a leading global provider of information and analytics, is using a custom, fine-tuned Claude 2 model to deliver conversational search, insightful summarization, and intelligent legal drafting capabilities via the company's new Lexis+ AI solution. Premier asset management firm **Bridgewater Associates** is developing an investment analyst assistant powered by Claude 2 to generate elaborate charts, compute financial indicators, and create summaries of the results. **Lonely Planet**, a renowned travel publisher, reduced its itinerary generation costs by almost 80 percent after deploying Claude 2; synthesizing its decades of travel content to deliver cohesive, highly accurate travel recommendations.

Anthropic and Amazon are both committed to the safe training and deployment of advanced foundation models. Amazon is an industry leader in cloud security and, as part of this agreement, is committed to promoting and implementing safety best practices on Amazon Bedrock to ensure the responsible use of our products and services. Both companies are actively engaged across a number of organizations to promote the responsible development and deployment of AI technologies, including the Global Partnership on AI (GPoA), the Partnership on AI (PAI), and the National Institute of Standards and Technology (NIST). Most recently, in July, both Amazon and Anthropic independently supported a set of voluntary safety commitments led by the White House to ensure that the future of transformative AI is guided by safety, security, and trust.

As part of the investment, Amazon will take a minority stake in Anthropic. Our corporate governance structure remains unchanged, with the [Long Term Benefit Trust](#) continuing to guide Anthropic in accordance with our [Responsible Scaling](#) principles. And this partnership will allow us to conduct deployment tests of new models to help us manage the risks.

of increasingly capable AI systems.

Training state-of-the-art models requires extensive resources including compute power and research programs. Amazon's investment and supply of AWS Trainium and Inferentia technology will ensure we're equipped to continue advancing the frontier of AI safety and research. We look forward to working closely with Amazon to responsibly scale adoption of Claude and deliver safe AI cloud technologies to organizations around the world.



## Related

[See All](#)

Company

**Thoughts on the US Executive Order, G7 Code of Conduct, and Bletchley Park Summit**

Nov 5, 2023 • 4 min read

Company

**Dario Amodei's prepared remarks from the AI Safety Summit on Anthropic's Responsible Scaling Policy**

Nov 1, 2023 • 5 min read

Company

**Anthropic's Responsible Scaling Policy**

Sep 19, 2023 • 4 min read

AI

Product

Research

Company

News

Careers

Press Inquiries

Support

Twitter

LinkedIn

Terms of Service

Privacy Policy

Acceptable Use Policy

Responsible Disclosure Policy

Compliance

© 2024 Anthropic PBC